

Securing California's Future: Artificial Intelligence

Daniel E. Ho¹
Stanford University

Testimony presented in the Hearing "California at the Forefront: Steering AI Towards Ethical Horizons" by the California Senate Governmental Organization Committee and the Senate Budget and Fiscal Review Subcommittee No. 4 on State Administration and General Government on February 21, 2024

Chair Padilla, Chair Dodd, and Members of the Committees, thanks for this opportunity to testify today.

My name is Daniel Ho and I'm a professor of law, political science, and computer science at Stanford University. I'm a senior fellow at Stanford's Institute for Human-Centered AI, serve on the National AI Advisory Committee, and I direct the Stanford RegLab, which works with a wide range of government agencies around AI demonstration projects.

California is the innovation capital of the world. As a nation, it would be the fifth largest economy. It houses 35 of the top 50 AI companies. Its educational institutions – from UC Berkeley to Stanford University to Caltech – are the envy of the world. Since World War II, the nation and California's partnership between government, universities, and private industry catalyzed fundamental advances that gave us microchips, GPS systems, and the worldwide web.

Yet this post-war model for innovation is under threat when it comes to AI.

One: AI research has become so capital-intensive that only a handful of private companies are at the frontier of the field. Vast amounts of computational power and data require huge capital investments. The innovation ecosystem has increasingly become more closed, concentrated, and opaque. That hurts science, accountability, and innovation.

¹ William Benjamin Scott and Luna M. Scott Professor of Law; Professor of Political Science; Professor of Computer Science (by courtesy); Senior Fellow, Stanford Institute for Human-Centered Artificial Intelligence (HAI); Senior Fellow at Stanford Institute for Economic Policy Research (SIEPR); Faculty Director, Stanford Regulation, Evaluation, and Governance Lab (RegLab), Stanford University. Member, National AI Advisory Committee (NAIAC), Administrative Conference of the United States (ACUS), Committee on National Statistics (CNSTAT). All views expressed in this testimony are provided in an individual capacity and do not represent the views of any affiliated institution, agency, or advisory committee. Thanks to Erika Bustamante, Isaac Cui, Lindsey Gailmard, Preeti Hehmeyer, Christie Lawrence, Olivia Martin, Anne Joseph O'Connell, Kit Rodolfa, Christine Tsang, Russell Wald, and Daniel Zhang for useful comments and conversations.

Two: Governance of this technology, as called for by the Governor’s executive order and much proposed legislation, is fundamentally challenging because government lacks AI expertise. Fewer than 1% of AI PhDs pursue public service. And government cannot govern AI if it does not understand AI.

Three: Many governmental systems still rely on “[dinosaur technology](#).”² Some 46M Americans turned to unemployment insurance during the pandemic, but twelve states, including California, still rely on dated software language from 1959.³ The system buckled, going from a timeliness rate of [97% to just above 50%](#).

So my core message to you is this: AI presents an extraordinary opportunity for the government and the people of California. But we must refrain and reject science fiction images of technology replacing humans: the best AI systems will support humans, not replace them.⁴

Let me give you one example. During the pandemic, RegLab worked extensively with Santa Clara County to use technology for Covid response. For example, we developed a system to help bilingual contact tracers connect with patients that needed [language assistance](#).⁵ That simple system improved same-day contact tracing by 12%, helping to reach those hit hardest early in the pandemic. The AI tool did not replace, but augmented the pandemic’s first responders. And [this kind of innovation](#) – which could modernize systems like unemployment insurance – is what we need in spades.⁶

However, realizing these benefits requires the responsible application of AI to these critical systems, and California must lead the nation by example for this responsible innovation. When developed poorly, AI systems have the potential to wreak significant harms: from amplifying biases to spewing misinformation to eroding privacy. Just as California has led on consumer privacy, California can help to rebalance our innovation system.

² Rebecca Rainey, *Democrats Call for UI System Fix as Millions Face Another Lapse in Benefits*, Politico (Feb. 10, 2021), <https://perma.cc/TWH6-D5XV>.

³ Makena Kelly, *Unemployment Checks Are Being Held Up by a Coding Language Almost Nobody Knows*, The Verge (Apr. 14, 2020), <https://perma.cc/VU2N-ZBAY>.

⁴ See, e.g. James Guszczka et al., *Hybrid Intelligence: A Paradigm for More Responsible Practice*, Center for Advanced Study in the Behavioral Sciences Working Paper (2022), <https://perma.cc/X3GZ-5NRP>; Erik Brynjolfsson, Danielle Li & Lindsey R. Raymond, *Generative AI at Work*, NBER Working Paper 31161, <https://www.nber.org/papers/w31161>; Nicholas Bien et al., *Deep-learning-assisted diagnosis for knee magnetic resonance imaging: Development and retrospective validation of MRNet*, 15 PLoS Medicine (2018), <https://journals.plos.org/plosmedicine/article/file?id=10.1371/journal.pmed.1002699&type=printable>.

⁵ Lisa Lu et al., *A Language-Matching Model to Improve Equity and Efficiency of COVID-19 Contact Tracing*, 118 PNAS e2109443118, <https://dho.stanford.edu/wp-content/uploads/e2109443118.pdf>.

⁶ Cal. Gov’t Operations Agency, State of California: Benefits and Risks of Generative Artificial Intelligence Report (Nov. 2023), <https://perma.cc/RM6L-K2EZ>.

Let me offer three recommendations of how we get there, about people, infrastructure, and regulation.

First, California must nurture, develop, and attract technical talent into public service. Most importantly, that means upskilling, namely providing opportunities for California's great civil servants to learn about the potential and risks of AI.

Similarly, California should consider models of the U.S. Digital Service, Digital Service Corps, and Presidential Innovation Fellows to hire AI talent.

California's great universities offer an extraordinary mechanism for upskilling civil servants and building a talent pipeline to the government. The Governor's GenAI Executive Order mandates that agencies consult with UC Berkeley and Stanford, and Stanford HAI is excited to co-sponsor California's AI Summit. But California should also follow the lead of the federal government, where agencies have built out academic partnerships to bring in expertise and develop pathways to public service for AI students.⁷ Senator Padilla's call for California Artificial Intelligence Research Hub, which could provide for a new model of public-private partnership, and a call for a [state talent exchange](#)⁸ to help upskill and augment the civil service are exactly the kind of initiatives that can be of central importance in helping the government seize this moment. California has one of, if not the, highest concentrations of AI talent in the world – California is in the perfect position to lead the country in building nimble programs to strengthen public sector tech expertise.

Second, a major reason why AI has become so concentrated in a small number of technology firms lies in technology infrastructure. The federal government has responded by piloting the [National AI Research Resource](#) (NAIRR),⁹ to democratize access to computing and data resources, based on a proposal and [White Paper](#) from Stanford HAI¹⁰ and a [federal task force](#).¹¹ The federal, bipartisan [CREATE AI Act](#) would fully scale the NAIRR.¹²

⁷ Daniel E. Ho, Opinion, *Gavin Newsom Signed an Executive Order on Artificial Intelligence. What Happens Now?*, Sac. Bee (Oct. 16, 2023); see generally Isaac Cui, Daniel E. Ho, Olivia Martin & Anne Joseph O'Connell, *Governing by Assignment*, 173 U. Pa. L. Rev. (forthcoming 2024), <https://dho.stanford.edu/wp-content/uploads/IPA.pdf>.

⁸ Daniel E. Ho, Anne Joseph O'Connell & Isaac Cui, Stan. Inst. Econ. Pol'y Rsch., *Talent Exchanges for State Governments* (Sept. 2023), <https://perma.cc/UUK8-CB7T>.

⁹ *National Artificial Intelligence Research Resource Pilot*, U.S. Nat'l Sci. Found., <https://perma.cc/3QVR-EL5N>.

¹⁰ Daniel E. Ho, Jennifer King, Russell C. Wald & Christopher Wen, Stan. U. Hum.-Centered Artificial Intelligence, *Building a National AI Research Source: A Blueprint for the National Research Cloud*, <https://perma.cc/M4MJ-9H2A>.

¹¹ *National Artificial Intelligence Research Resource Task Force Releases Final Report*, White House Off. Sci. & Tech. Pol'y (Jan. 24, 2023), <https://perma.cc/RB8E-5LWZ>.

¹² U.S. Senate Artificial Intelligence Caucus, *Create AI Act of 2023*, <https://perma.cc/TJG4-RKPC>; S. 2714, 118th Cong. (2023).

California should lead the nation in these efforts. Proposals for computing and data resources like CalCompute or the AI Research Hub could broaden access to a wider range of Californians that are increasingly left out. Government has a major advantage to re-align technology with human values: reliability of data. Currently, AI is powered by hoovering up all data on the internet and learning from it. But garbage in, garbage out. Toxicity, falsehoods, and risks from GenAI are a result of this choice. By providing secure, privacy-protected access to much higher quality government data, the NAIRR and California proposals can turn AI to solve much more socially useful problems.

One example is that when the U.S. Geological Service made (Landsat) satellite imagery free, it created [3 to 4 billion dollars](#) in value annually and drove forward our ability to understand and respond to environmental threats.¹³

The state that provides this infrastructure will lead in AI.

Third, when it comes to [governing AI](#),¹⁴ California must address the huge information asymmetry about new and heightened risks posed by AI. The most urgent and lower-budget item [regulatory solution](#)¹⁵ hence lies in [adverse event reporting](#).¹⁶ Just as in cybersecurity, where parties are required to report vulnerabilities and attacks, we need such a system for AI, enabling government to monitor, investigate, and respond to emergent risks of AI systems.

It's a problem when only a small number of self-interested actors have the information necessary to describe risks. Over the past year, for example, there was lots of worry about how ChatGPT-like systems could help hostile actors create bioweapons. But the evidence was paper thin and has been debunked. [RAND](#), for instance, showed¹⁷ that such systems (a) offer no information beyond what is readily available on the internet, and (b) yield no (statistically significant) difference in the ability to create a bio-attack plan.

Government must be able to separate hype from reality. The way we've addressed this in other areas – cybersecurity, drugs and medical devices, pathogens, and design defects – is adverse event reporting. That's what's needed for AI.

¹³ Crista L. Straub, Stephen R. Koontz & John B. Loomis, U.S. Geological Survey, Open-File Report 2019-1112, Economic Valuation of Landsat Imagery (2019), <https://perma.cc/YC5N-35WL>.

¹⁴ Neel Guha et al., *AI Regulation Has Its Own Alignment Problem: The Technical and Institutional Feasibility of Disclosure, Registration, Licensing, and Auditing*, 92 Geo. Wash. L. Rev. (forthcoming 2024), https://dho.stanford.edu/wp-content/uploads/AI_Regulation.pdf.

¹⁵ Neel Guha et al., Stan. U. Hum.-Centered Artificial Intelligence, The AI Regulatory Alignment Problem (Nov. 2023), <https://perma.cc/9J5J-PYEH>.

¹⁶ Nat'l Artificial Intelligence Advisory Comm., RECOMMENDATION: Improve Monitoring of Emerging Risks from AI through Adverse Event Reporting (Nov. 2023), <https://perma.cc/8VE7-2XHZ>.

¹⁷ Christopher A. Mouton, Caleb Lucas & Ella Guest, RAND, The Operational Risks of AI in Large-Scale Biological Attacks: Results of a Red-Team Study (Jan. 2024), <https://perma.cc/945G-LHK8>.

Government procurement of AI tools and systems can also be a powerful lever for increasing transparency and other trustworthy AI practices. California cities like [San Jose](#) are leading efforts through a [national AI coalition](#) to standardize procurement and make sure vendors provide the government with information necessary to justify public trust in the government's use of AI.

Let me conclude with two last words on AI regulation. One is that we must be extremely careful about regulation that has the potential to entrench incumbents and quash [open innovation](#).¹⁸ The poster child of that kind of ill-conceived regulation is a [licensing regime](#),¹⁹ whereby only a few, well-equipped companies would get licenses to develop advanced AI models. This is wrong and would reduce oversight and stifle innovation.

Another is that many regulatory proposals single out AI systems as being uniquely risky and subject to controls. But AI can also expose vulnerabilities of existing human systems. In developing AI systems for the IRS, for instance, our team uncovered disturbing [racial disparities in tax audits](#)²⁰ of existing legacy systems (not fancy AI). While many worry about AI and biorisk, the real vulnerability lies as much in the existing practice of allowing laboratories to self-regulate the sale of [synthetic DNA sequences](#).²¹ Concerns about AI may actually point us to opportunities for reform of existing systems.

In sum, AI has tremendous promise. But just as post-war California set us on the path to become the innovation engine of the world, we must take these steps now to ensure that we retain this leadership both for technology and our values.

I welcome your questions.

¹⁸ Rishi Bommasani et al., Stan. U. Hum.-Centered Artificial Intelligence, Considerations for Governing Open Foundation Models (Dec. 2023), <https://perma.cc/V5Y8-QU9G>.

¹⁹ Guha et al., *AI Regulation Has Its Own Alignment Problem*, *supra* note 14.

²⁰ Hadi Elzayn et al., Stan. Inst. Econ. & Pol'y Rsch., Measuring and Mitigating Racial Disparities in Tax Audits (2023), https://dho.stanford.edu/wp-content/uploads/IRS_Disparities.pdf.

²¹ Nat'l Acads., Biodefense in the Age of Synthetic Biology (2018), <https://perma.cc/U5RF-7DVR>.